PCIe® Cable Update

Alex Haser
OCuLink WG Chair
Molex, LLC

Disclaimer

Some of the information in this presentation refers to specifications still in the development process. This presentation reflects the current thinking of various PCI-SIG® workgroups, but all material is subject to change before the specifications are released.

Agenda:

- **ECNs & Errata Completed Against OCuLink 1.0:**
  - Memory Map ECN
  - Skew ECN
  - Server Space ECN
  - Pin Sequencing Errata
- **Open Items:**
  - BP Type ECR
  - CPRSNT# ECR
  - Wiring Chart ECR
  - Performance Table ECR
  - 16.0 GT/s Performance ECR

ECNs & Errata Completed Against OCuLink 1.0

- **Memory Map ECN :**
  - Modifies the cable assembly memory map; reorganizes bytes for external cable assemblies and modifies fields as needed to align with SFF-8636
  - Impact: Existing external OCuLink cables and ports must be modified to accept the new memory map
  - Completed: 12/13/2016
- **Skew ECN:**
  - Corrects maximum skew permitted in OCuLink cables to coincide with requirements in *CEM* specification
  - Impact: Minimal; manufacture of compliant cable assemblies may be more challenging
  - Completed: 12/13/2016

ECNs & Errata Completed Against OCuLink 1.0

- **Server Space ECN:**
  - Adds environmental requirements for enterprise applications to *OCuLink 1.0* specification
  - Impact: Minimal environmental & mechanical requirements for OCuLink connectors and cables are identified for use in enterprise applications
  - Completed: 12/13/2016
- **Pin Sequencing Errata:**
  - Contains one pin sequencing erratum against PCIe *OCuLink 1.0*, correcting a single entry in a table showing the pin sequencing for a ground pin as second mate
  - Completed: 11/04/2016

Open Items:

- **BP Type ECR:**
  - Fully defines the Backplane (BP) Type signal that was incompletely defined in the *OCuLink 1.0* release
  - Impact: BP Type pin may be used for VSP after BP Type is established
  - Expected to be completed: early Q3 2017 (in review)
- **CPRSNT# ECR:**
  - Fully defines the cable presence (CPRNST#) signal that was incompletely and inaccurately defined in *OCuLink 1.0*
  - Impact: CPRNST# supported 3 states in *OCuLink 1.0*; it now only supports 2
  - Expected to be completed: Q2 2017 (in review)

Open Items:

- **Wiring Chart ECR:**
  - Connector and cable assembly pinout tables have been revised to show complete pinout assignments for all applications
  - Impact: Less chance of errors
  - Expected to be completed: Q3 2017 (in review)
- **x4 Figure clean-up ECR:**
  - Replaced drawings showing SMT hold-downs with drawings showing through-hole hold downs to reflect the industry's preference
  - Corrects dimensions and tolerances based on industry feedback
  - Impact: Accurate and concise representation of x4 form factor
  - Expected to be completes: Q4 2017

Open Items (Continued):

- **Performance tables ECR:**
  - Reorganizes and clarifies mechanical and environmental performance criteria required of OCuLink connectors and cables
  - Impact: Test requirements adhere to EIA test methodology
  - Expected to be completed: Q4 2017
- **16.0 GT/s Performance ECR**
  - Adds signal integrity performance requirements necessary to allow OCuLink cables to meet 16GT/s system requirements
  - Impact: OCuLink cables may be used in 16GT/s systems
  - Expected to be completed: Q1 2018

PCIe® Cable Update

Lee Mohrmann
External Cabling WG Chair
National Instruments

Disclaimer

The information in this presentation refers to specifications still in the development process. This presentation reflects the current thinking of various PCI-SIG® workgroups, but all material is subject to change before the specifications are released.

Current Cabling Status

- **1.0 Specification released Jan 2007**
- **2.0 Specification released Q3 2012**
- **3.0 Specification 0.7 draft completed Q2 2016**
  - 0.9 draft completed Q2'17

PCI Express® External Cable Connectors for 3.0 and Beyond

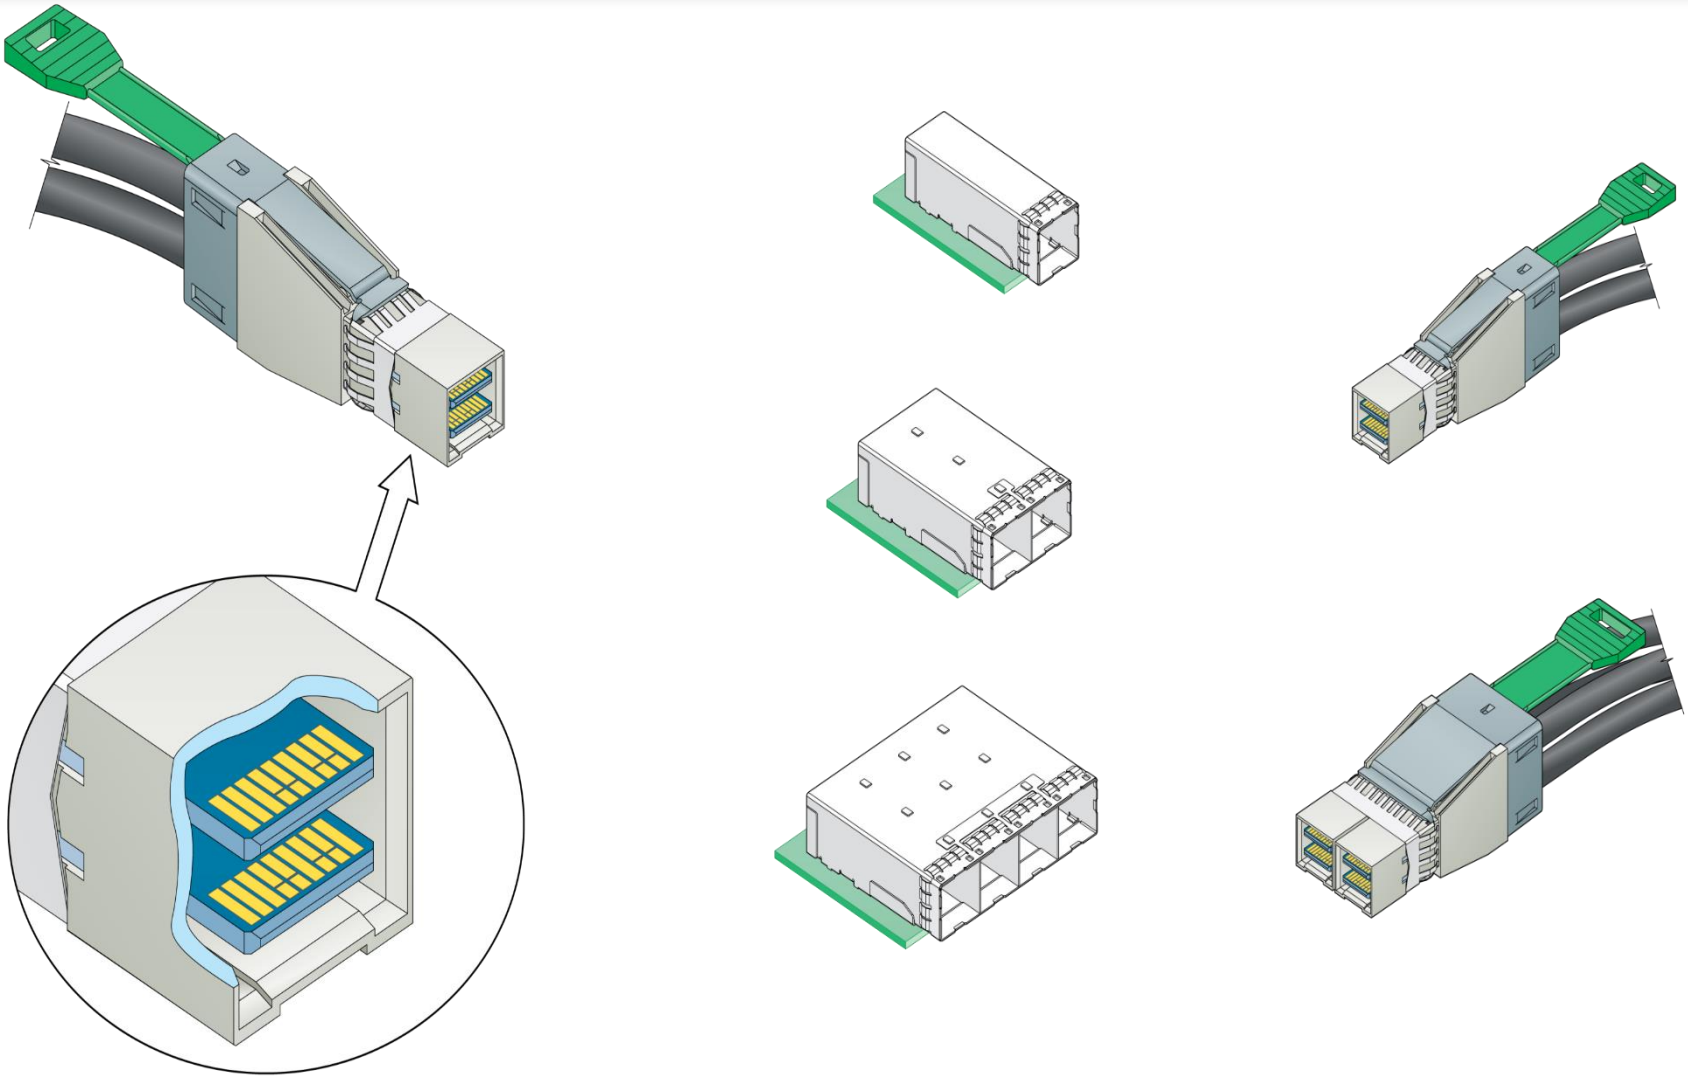

Cabled System Hierarchy

In general, PCIe® assumes common power supply and reset timings.

The Subsystem terminology in ExtCable spec 2.0 only specifies what's attached to each end of the cable.

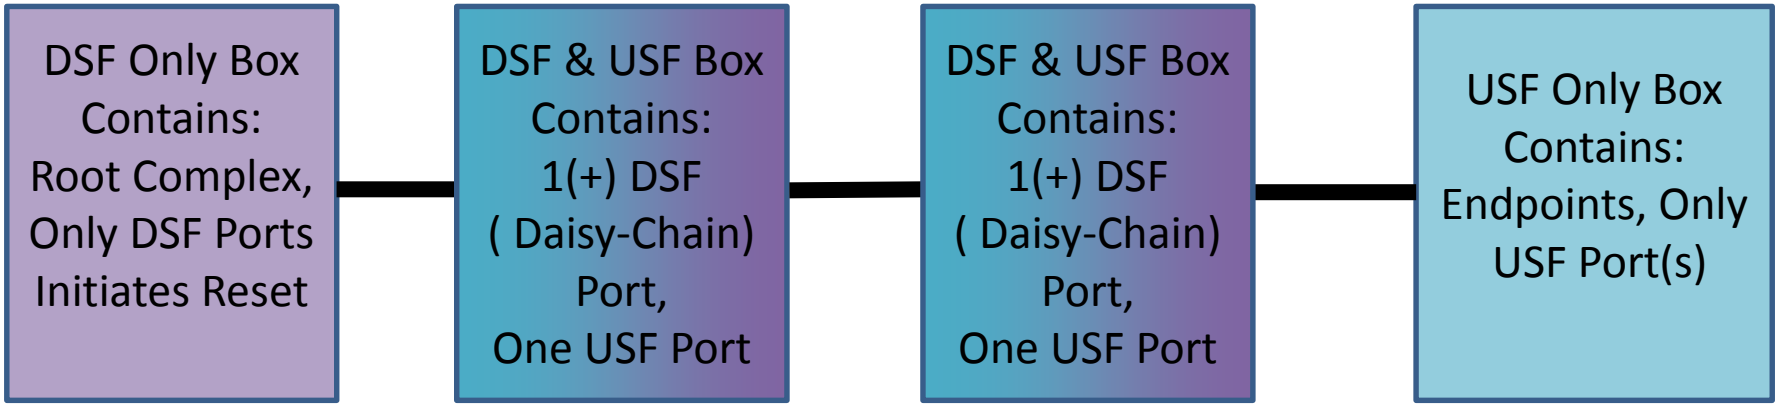

For the 3.0 specification, the concept of a large expansion system in included. Terminology used to reference specific locations.

Terminology

Location Term	Description
Cable Link Partner, CLP	The Fixed-Side Subsystem that a device and port are connected to via a cable
Downstream Facing Port, DSF	A cabled port in the Upstream Subsystem that will ink over the cable to a Downstream Subsystem.
Local	A Fixed-Side port and inserted cable connector.
Upstream Facing Port, USF	A cabled port in the Downstream Subsystem that will link over the cable to an Upstream Subsystem.

Terminology Illustration

Reference *Box*,
Local describes any signals unique to this *Box*

Cabled Link Partner (CLP)
with respect to reference box

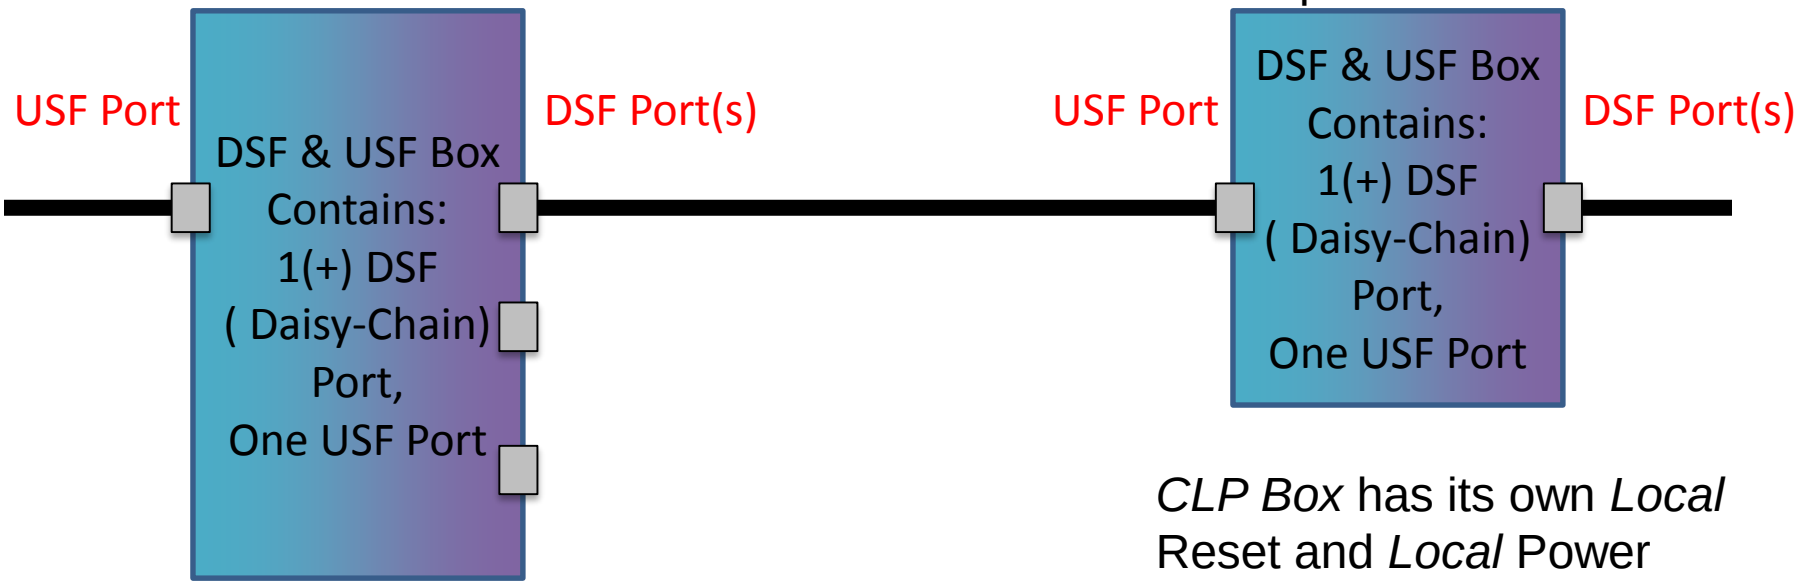

Box has a *Local* Reset, a *Local* Power Good

CLP Box has its own *Local* Reset and *Local* Power Good

Cable Management Interface

- **Attaching an external device to a PCIe 8GT/s interface has some challenges.**
  - Link training without preset hints based on channel knowledge is out of the norm in the PCIe ecosystem.

- **The new specification includes a 2-wire interface that:**
  - Enables Configuration of the Links
  - Performs Sideband Functions
  - Enables Cable Manageability
  - Allows for Vendor Specific Messages OOB
    - Future architectures for OOB messages for PCISIG form factors fits into existing VSM architecture, at a much slower rate

- **This was determined to be the best way to transfer the cable assembly information to the PCIe silicon for SERDES and port configuration.**

Structure of the 2-wire Interface

- **Management busses, such as IC2 and SMBus, have existed for some time. Most developers have run into concerns regarding too many device loads, odd topologies, accelerators, level shifters, etc., that create issues on the bus.**
- **The specification for the Cable Management Interface will allow only four devices per interface.**
  - Controllers that support port bifurcation will need multiple bus controllers to handle the multiple links created. A separate interface will be required to interface with the PCIe device or any other system interface.
  - This is not a specification issue for external devices, it is anticipated that a downstream cable port may only connect to a single upstream port.

Cable Management Interface Controller for Upstream Port

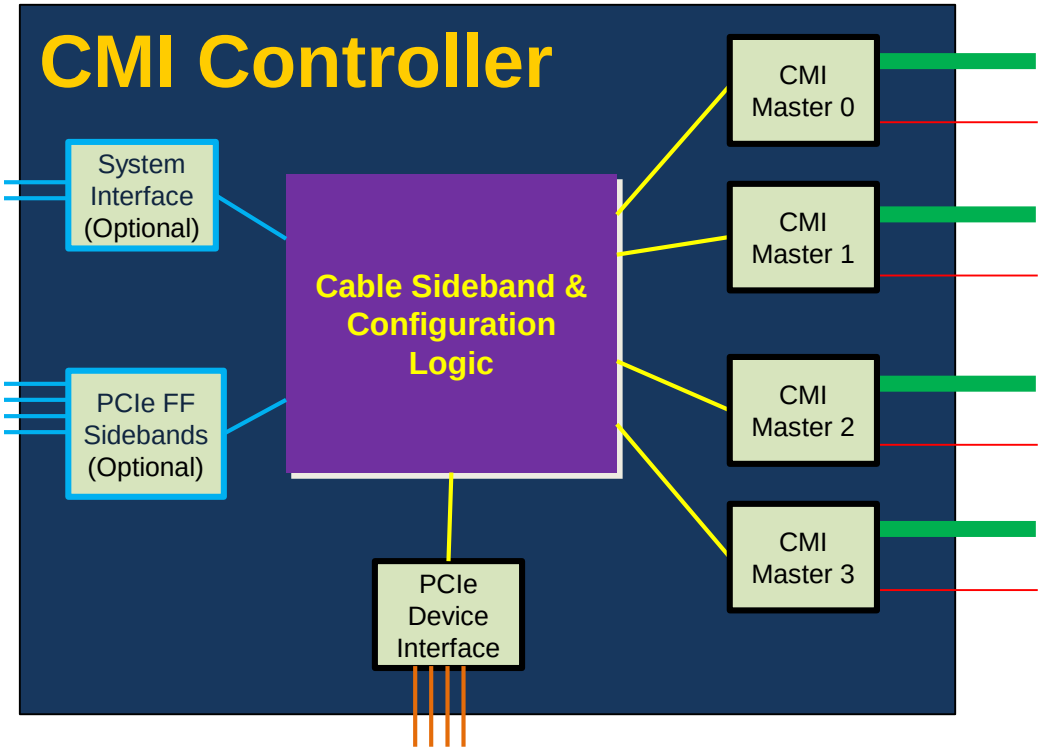

This Cable Management Interface (and sideband) Controller must interface each individually supported cable interface and the PCIe device.
The controller has the option to interface to PCIe sidebands for the form-factor where the controller resides, and to an external system management interface.

CMI Configurations

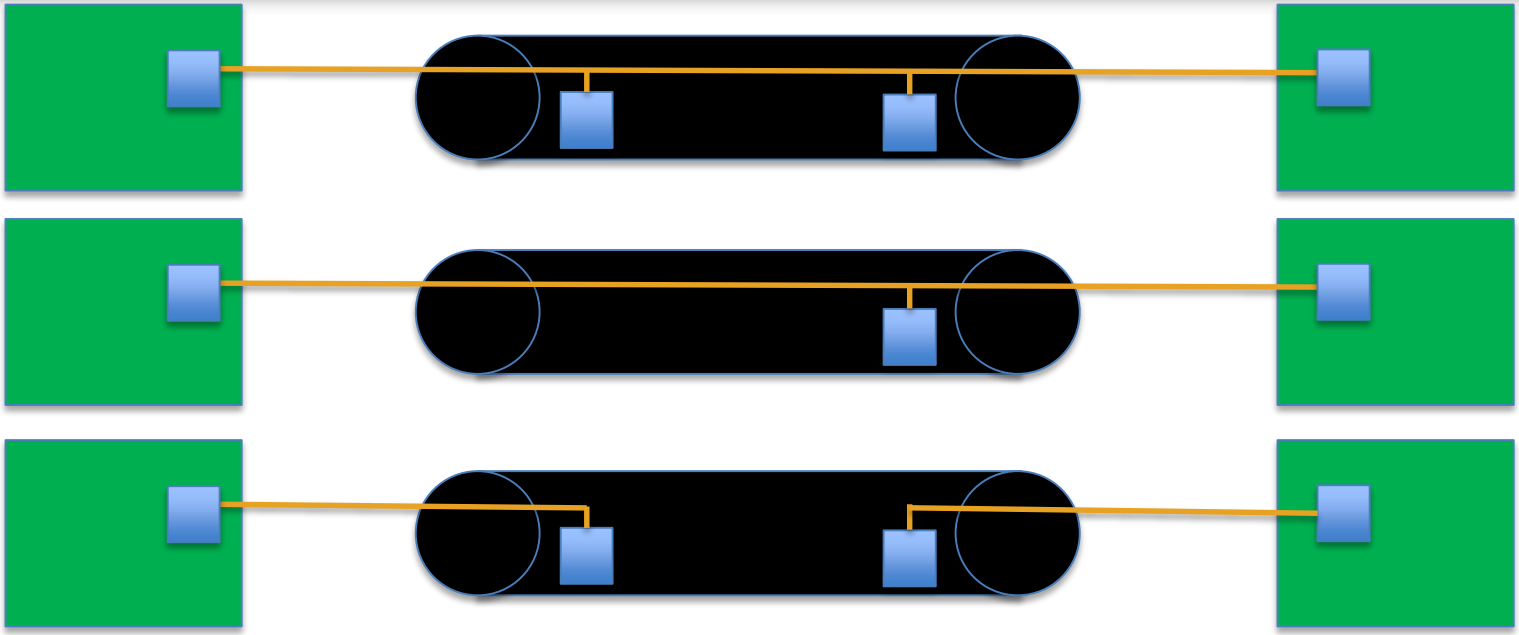

There are three fundamental configurations.

- **Sideband-Enabled, Dual Cable Devices**
- **Sideband-Enabled, Single Cable Device**
- **No Cable Sideband Messages, Dual Cable Device**

CMI Functional Corner Cases

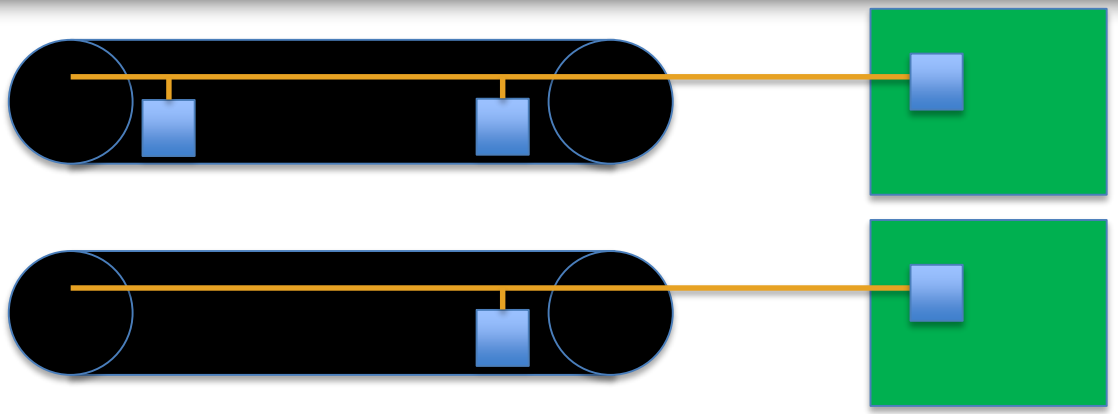

- **Sideband-Enabled cable assemblies can present challenges for the CMI signaling when only one end is attached and powered.**
  - No power lines traverses the cable, only one end is powered via the PWR and MGTPWR pins.
  - Signal integrity and timing can be problematic over a run with a long stub, even at 100 kHz.
  - Only one pullup resistor is present in this scenario.

Challenges for the CMI Bus

- **CMI Controller must be able to communicate with the near end cable device when powered, regardless of status of the other end.**
  - One of the things that must be determined is if the cable supports sideband messages.
- **For 5m cables, this could be a very long stub.**
- **The CMI controllers and cable devices have their own capacitances. Is there a limit?**
- **What's the impact to the signal if there's no impedance control?**
  - Impact of impedance on capacitance?
  - Impact of capacitance on length?
- **And we want to be compatible with SFF-8449 cable assemblies…**

CMI Signal Integrity Example

We had to balance the ability to manufacture assemblies with higher impedance wires with the signal integrity issues that could occur with low impedance wires.

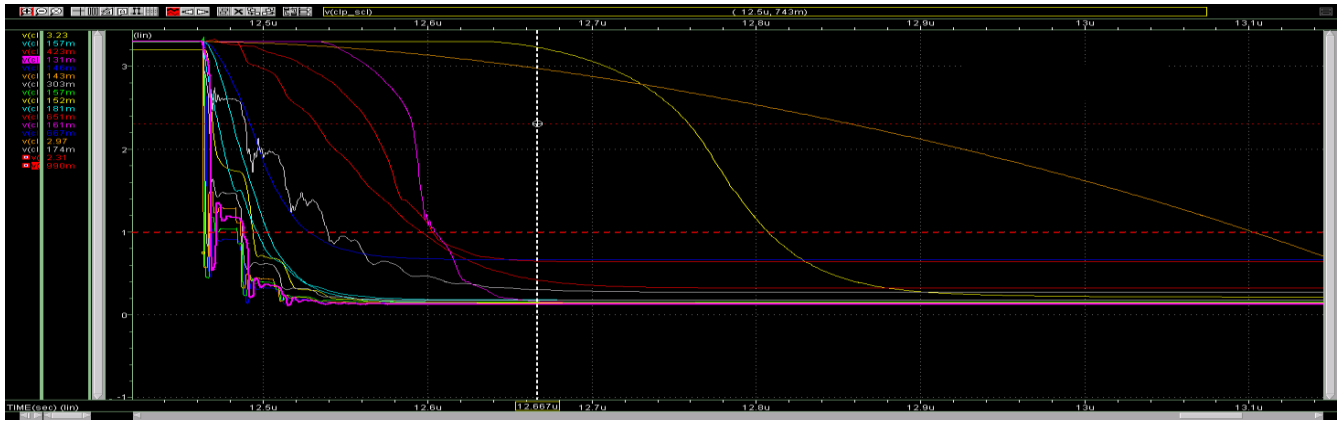

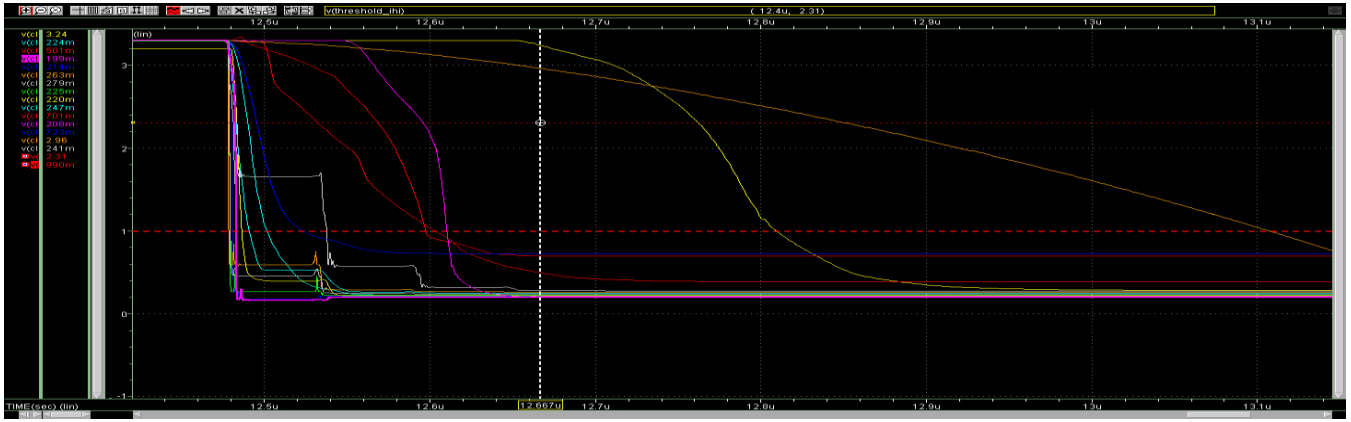

Other CMI Obstacles

- **Pullup selection must meet both signal integrity and loading**
  - A strong pullup resistor enables a quick T_{HIGH} to deal with highly capacitive loads due to devices and channel.
  - A weak pullup allows a good V_{LOW} level with a wide variety of devices.

- **Cable Impedance must meet requirements for capacitive load, length, and manufacturability.**

Snapshot of Analysis

- 2.2k is closest 5% resistor to 2.13 kΩ needed for 3mA driver to reach VOL of 0.4V
- Setup implements 2.2k+5% to provide weakest pullup for Single-Side connection to cable.

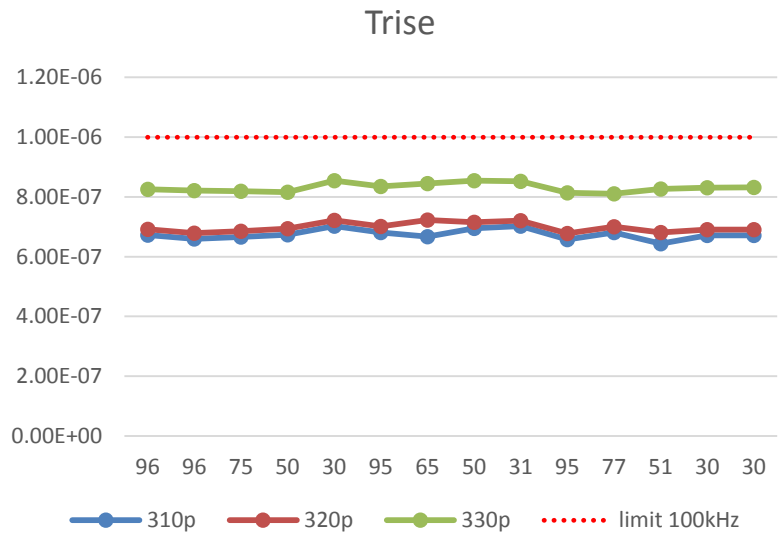

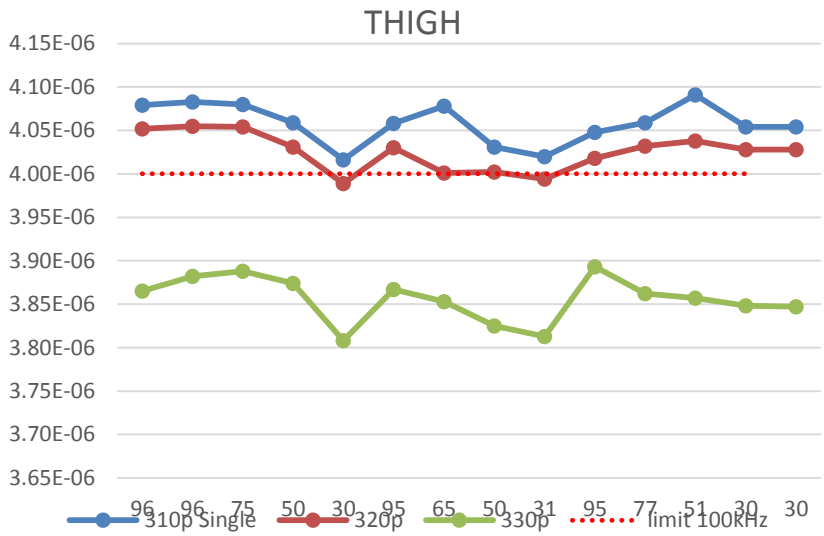

The Legends indicate the total wire capacitance of the cable

Summary of CMI Implementation Details for 0.7 Draft.

- **CMI Pullups of 2.15k on the Fixed-side sub-system(s)**
- **I_{SINK} of 3.0 mA for CMI devices**
- **Impedance target of PCB traces for CMI is 55 Ω**
- **Impedance of wires for CMI is >50 Ω**
- **Total Capacitance of Cable Assembly <330 pF**
- **100 kHz operation only***

***Future work for ECN or 4.0 to include 400kHz operation.**

Enabling Power to Large Systems

- **Cabled PCIe Implementations differ from “in the box” form factors due to their independent power systems.**
- **CMI_Power_Enable initiates power in the downstream direction**
- **CMI_Sstart allows power to be remotely enabled in a more deliberate fashion**
  - There are two types of Sstart that further allow staggering of powerup: gated DSF and Local First.
- **CMI_Wake initiates power in the upstream direction**

Cabled System Power

Each Box below is a system with independent power supplies and local reset structures.

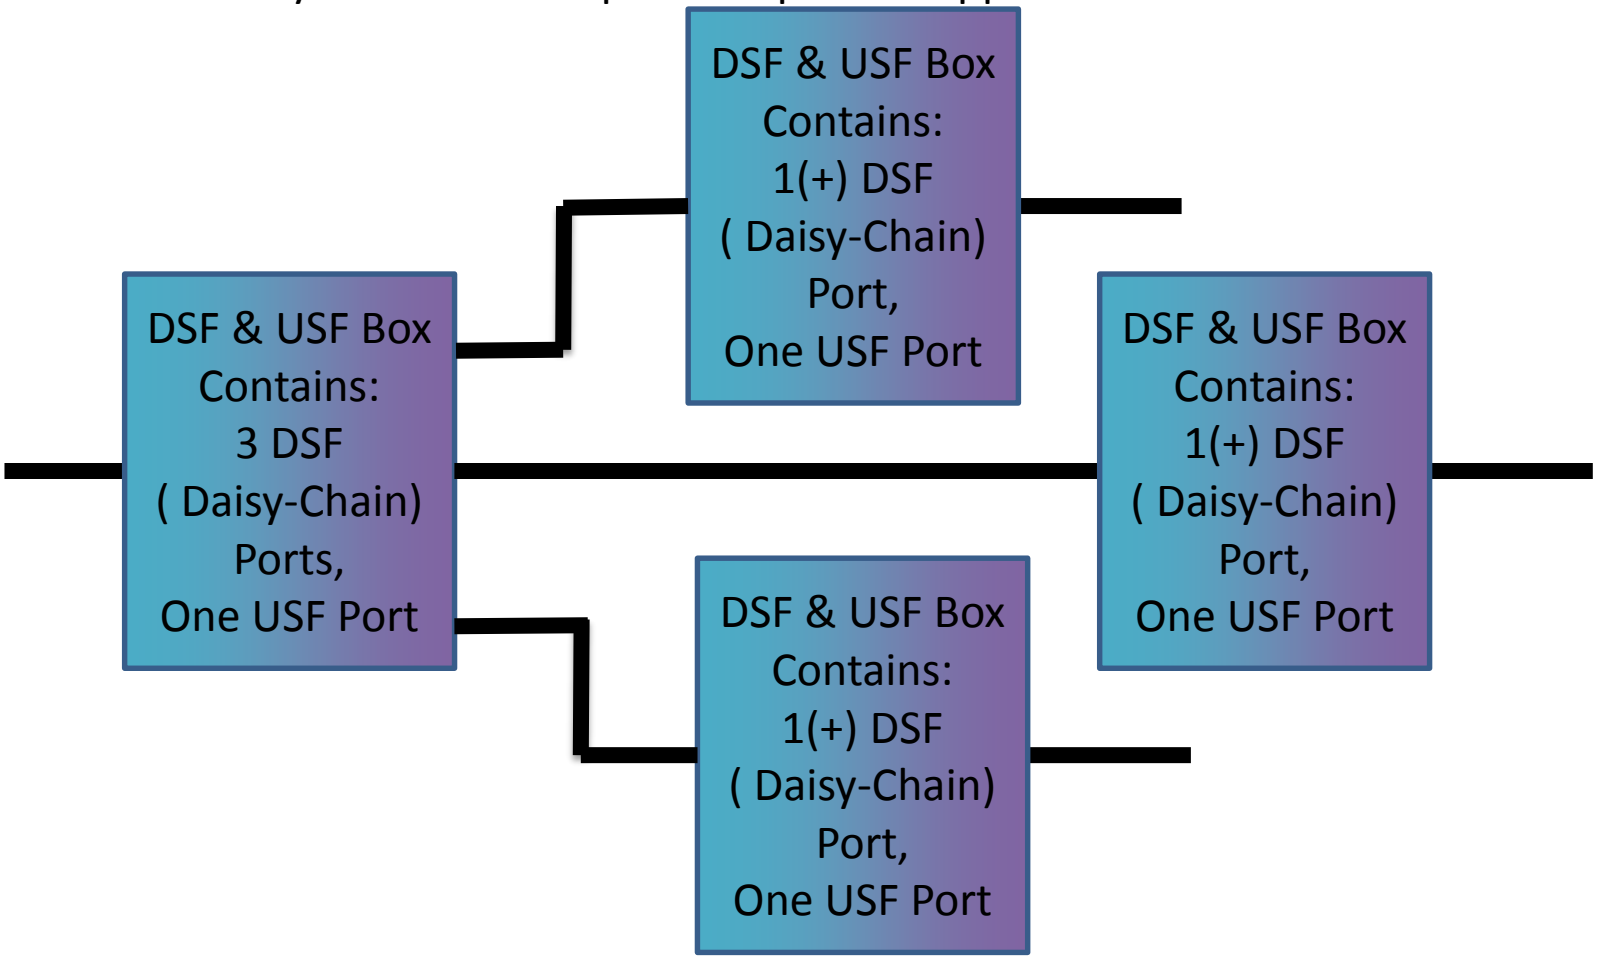

Staggered Start Implementations

- **Local First**
  - Powers self before sending CMI_POWER_ENABLE and CMI_SSTART to downstream ports. The commands may be staggered per downstream port for further separation of power initiation.
- **DSF Gated**
  - Sends CMI_POWER_ENABLE and CMI_SSTART to downstream ports before enabling local power. The commands may be staggered per downstream port for further separation of power initiation.
- **In either case, CMI_Wake and CMI_SSTART are not propagated to the USF port until messages are received with CMI_CLP_READY on all DSF ports.***

***These are ignored for ports with cables that do not support Sideband Messages.**

DSF-Gated Power Sequence

Each Box below is a system with independent power supplies and local reset structures.

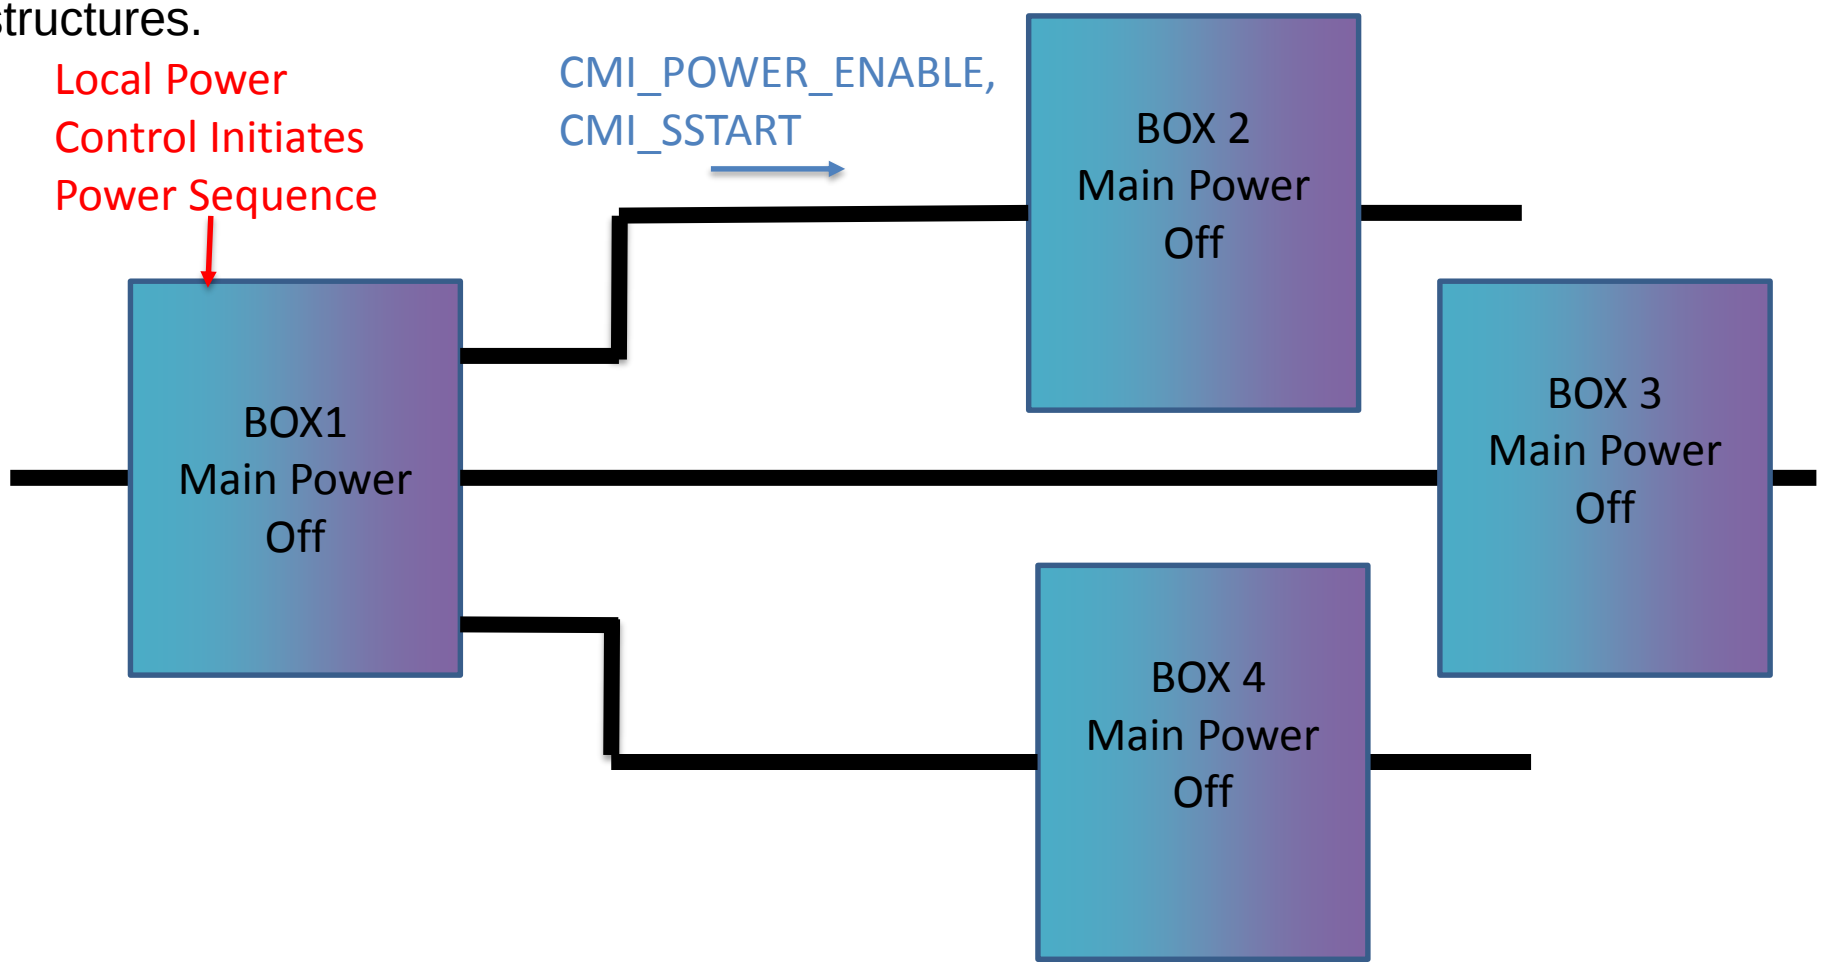

DSF Gated Power Sequence

Each Box below is a system with independent power supplies and local reset structures.

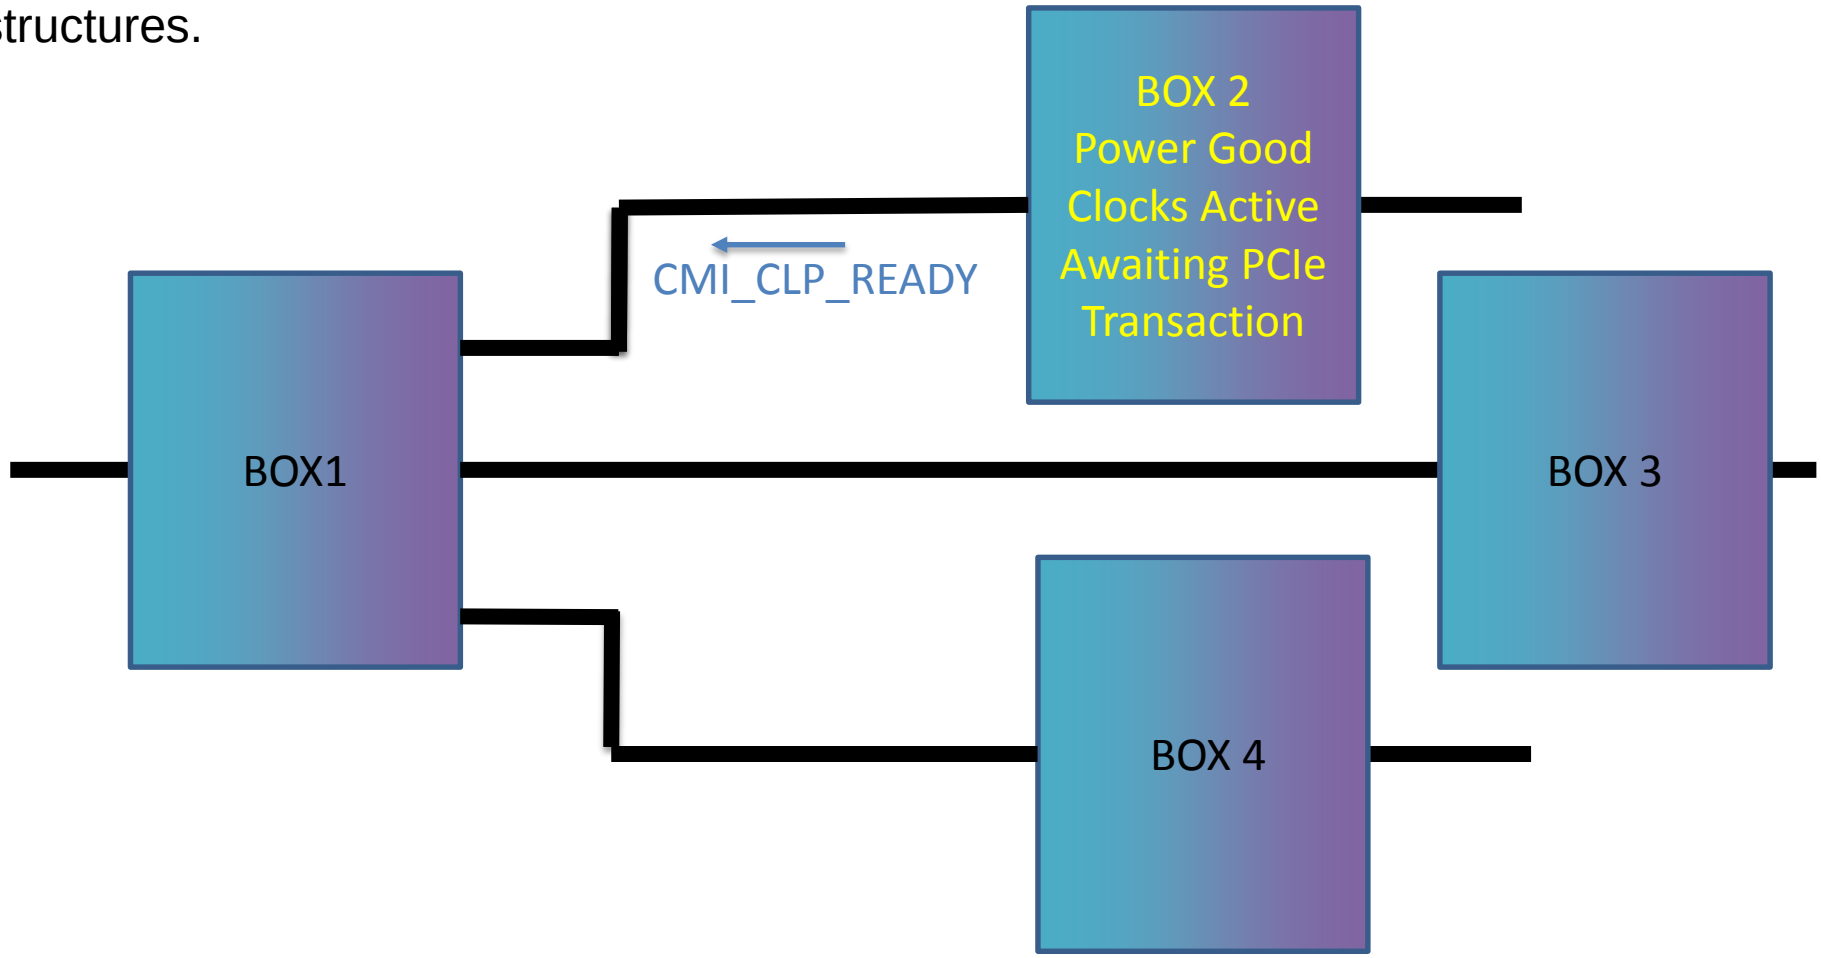

DSF Gated Power Sequence

Each Box below is a system with independent power supplies and local reset structures.

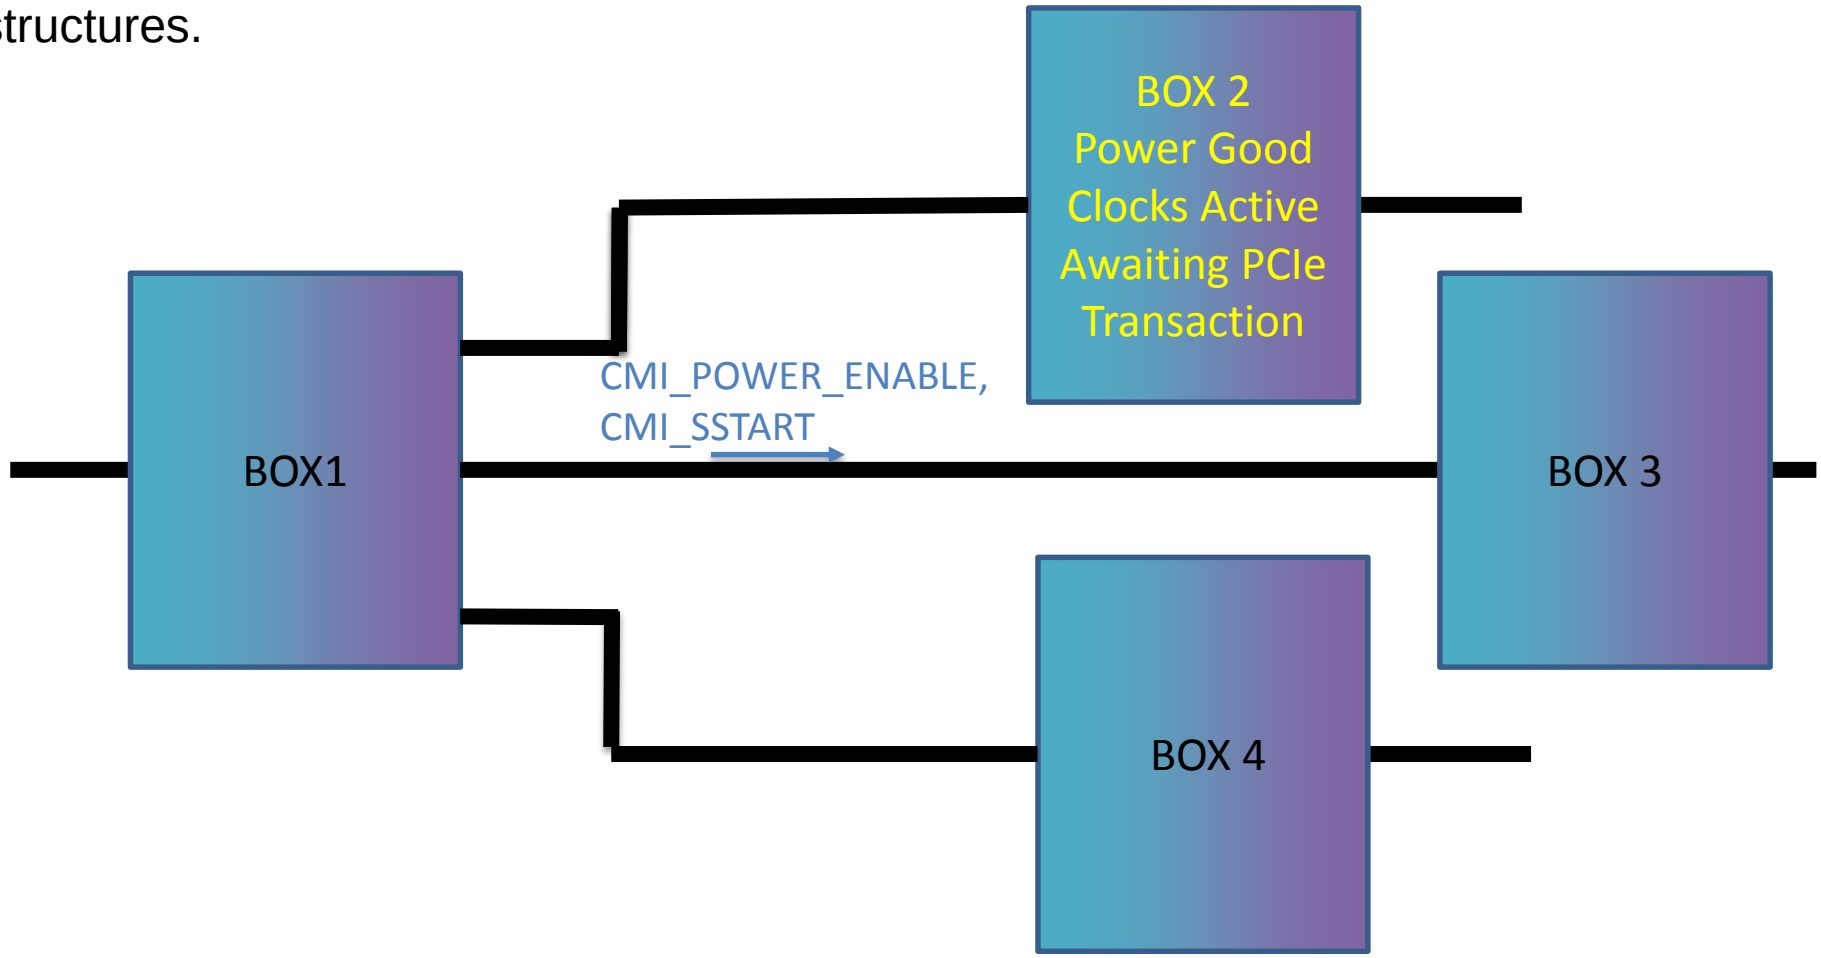

DSF Gated Power Sequence

Each Box below is a system with independent power supplies and local reset structures.

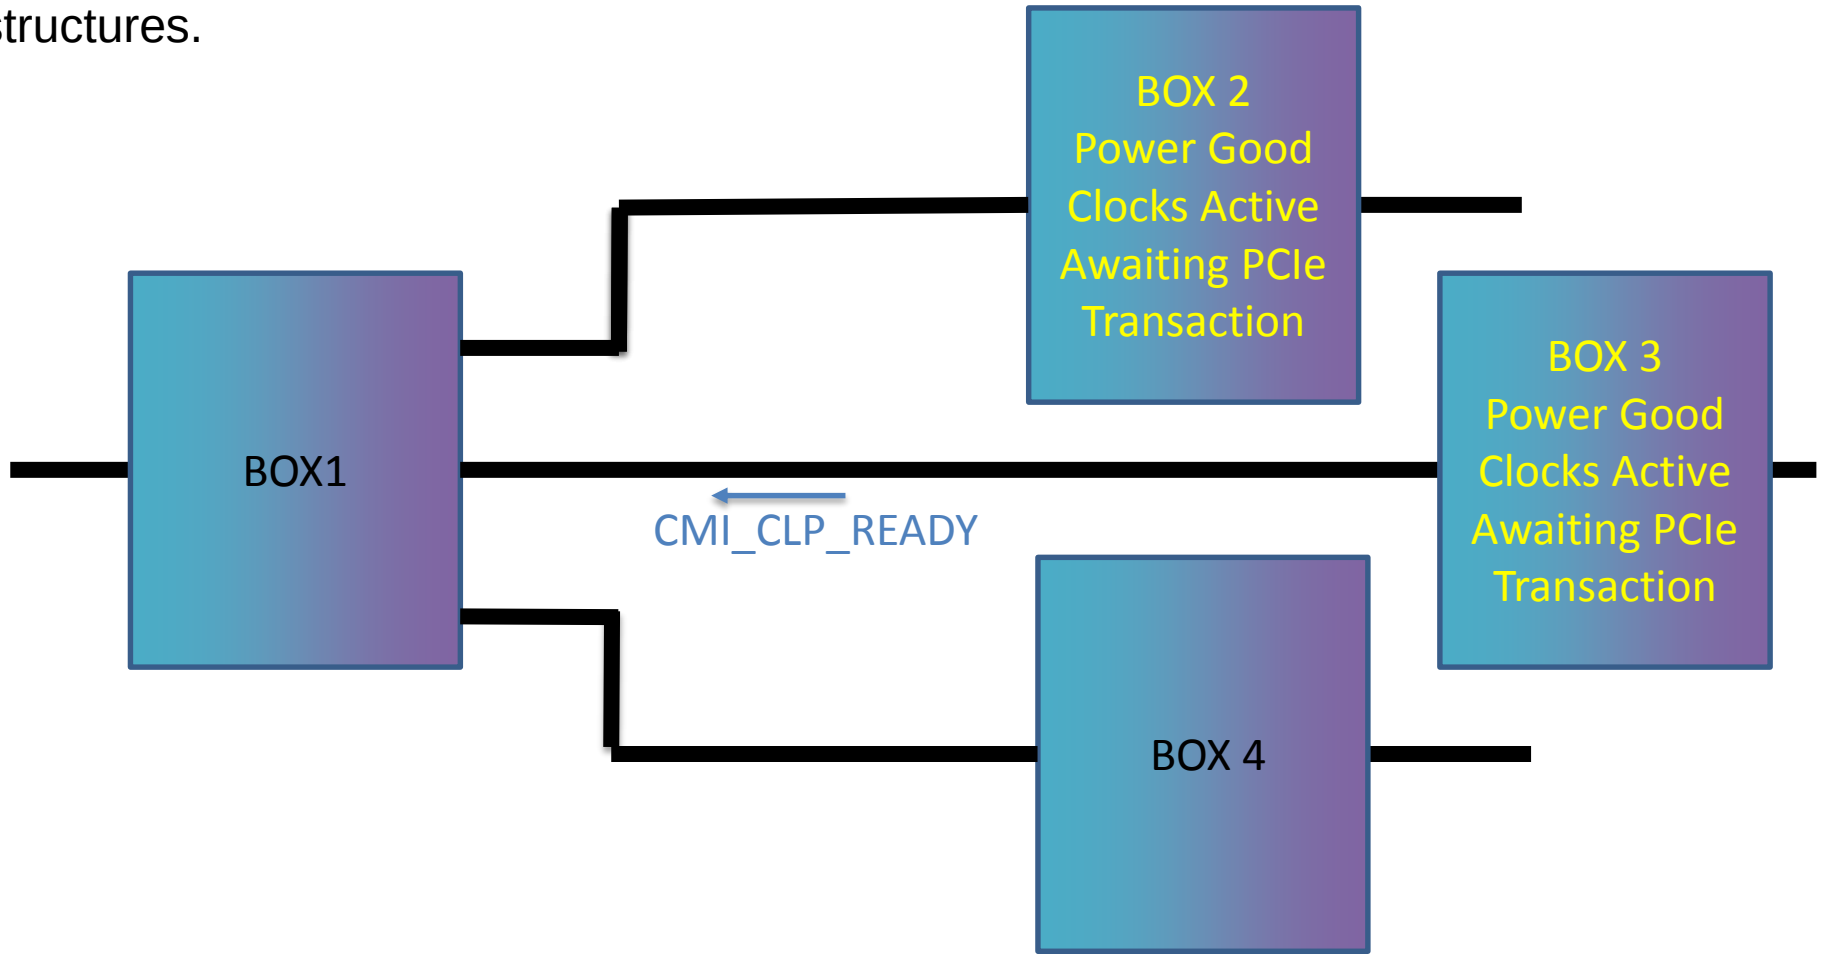

DSF Gated Power Sequence

Each Box below is a system with independent power supplies and local reset structures.

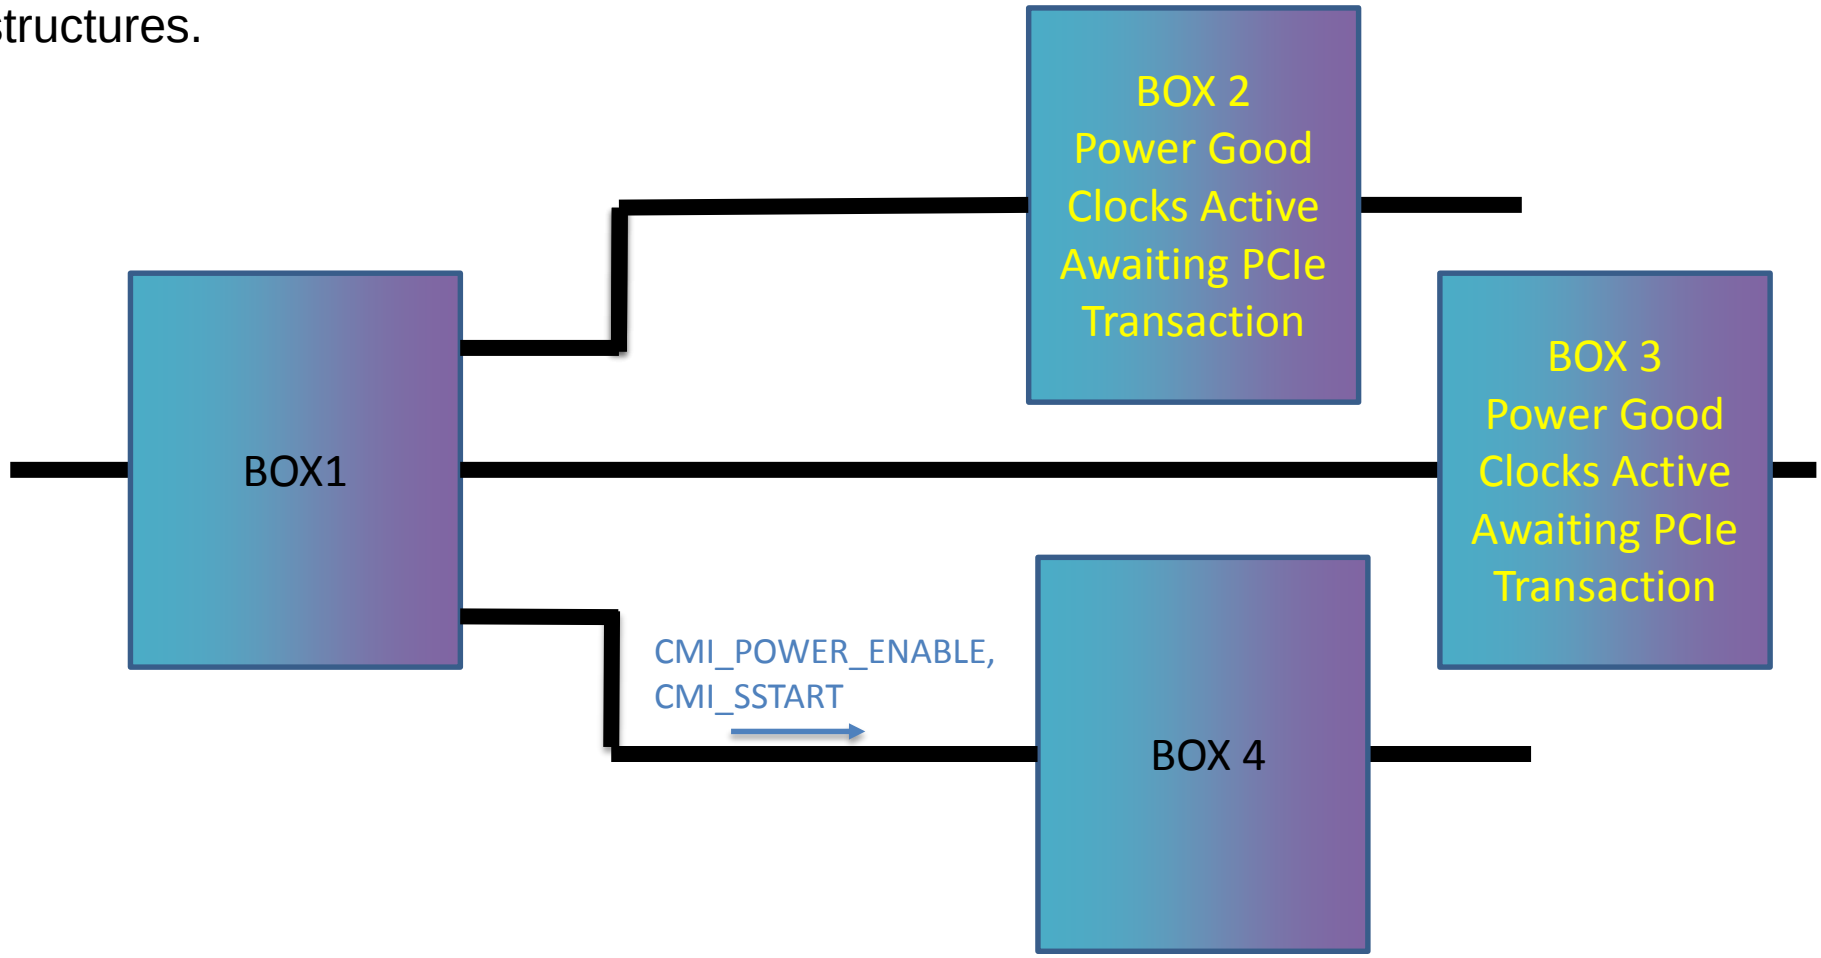

DSF Gated Power Sequence

Each Box below is a system with independent power supplies and local reset structures.

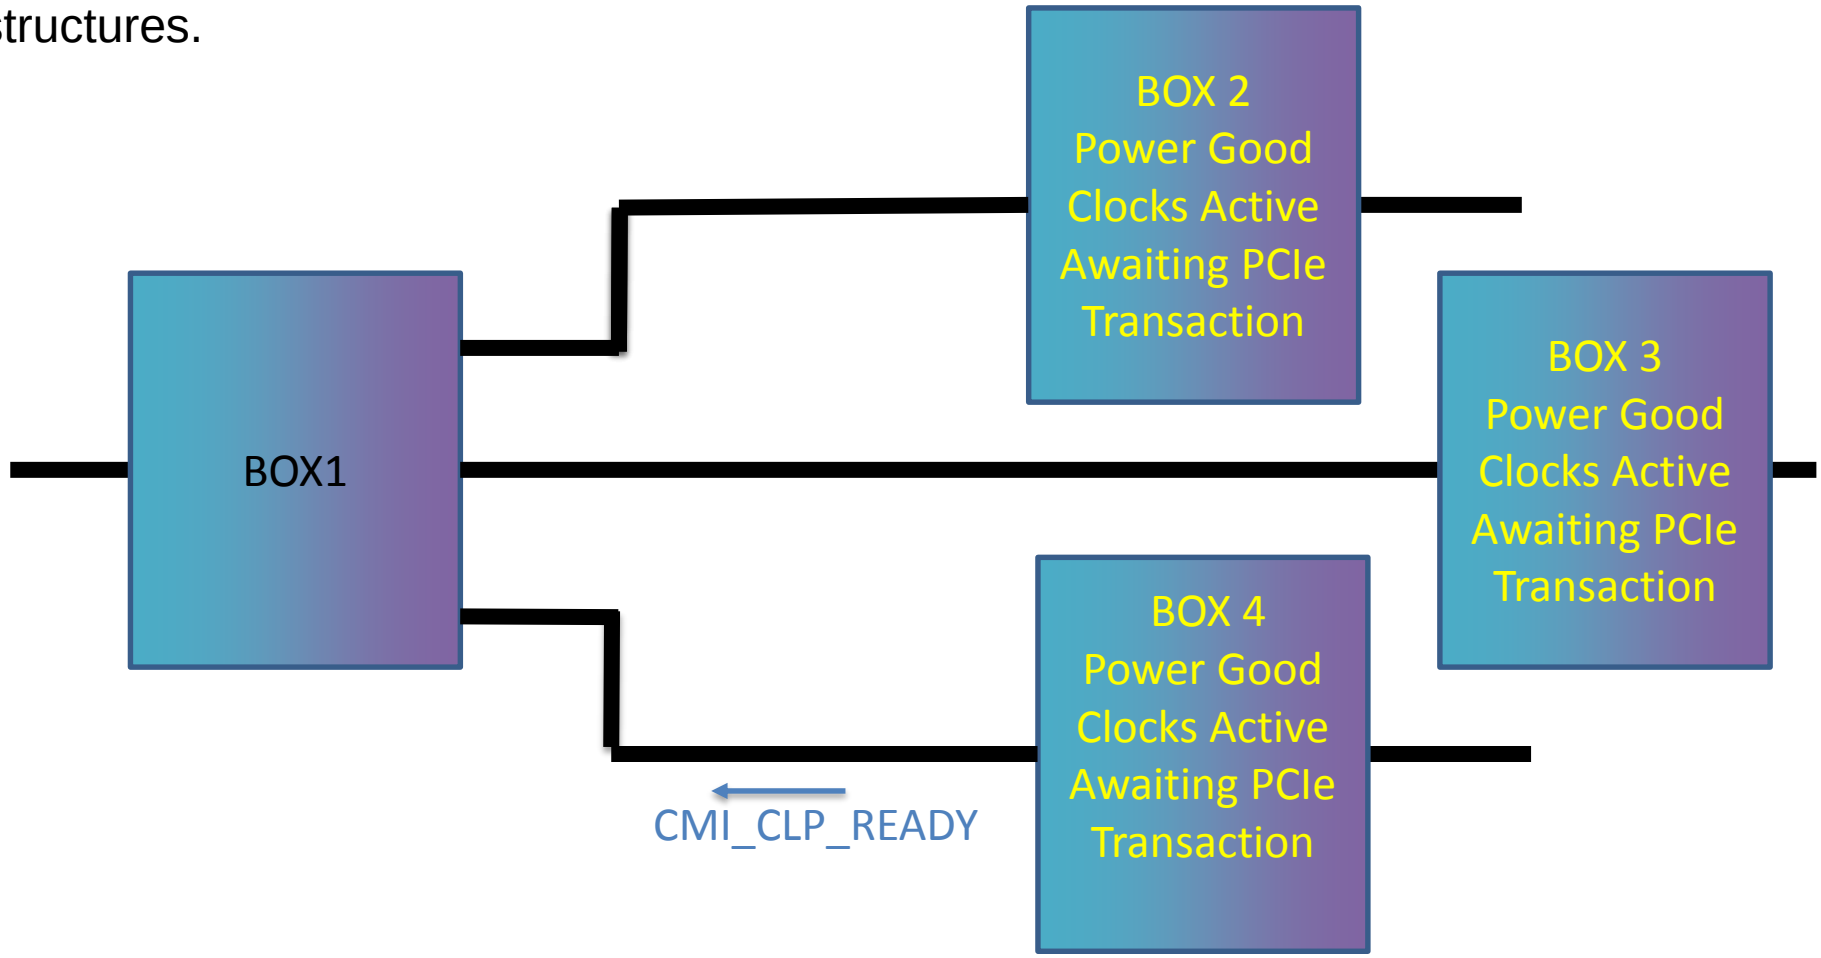

DSF Gated Power Sequence

Each Box below is a system with independent power supplies and local reset structures.

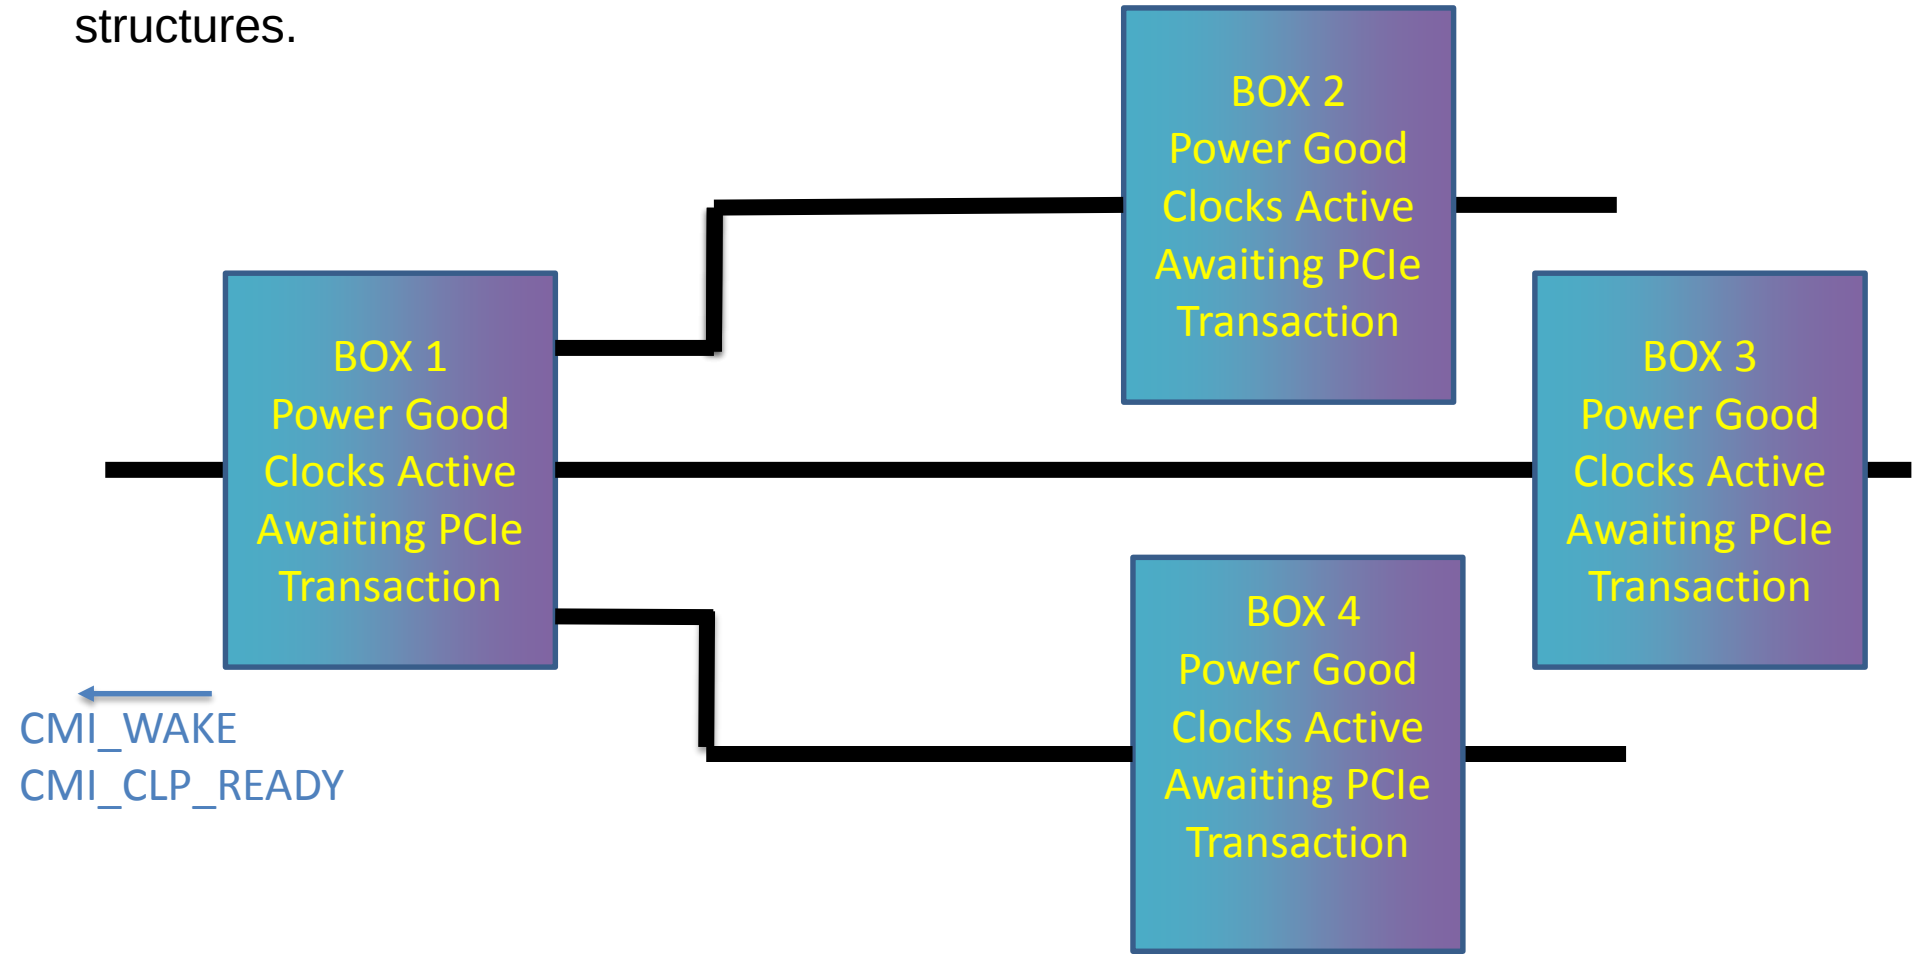

Local First Gated Power Sequence

- **This is similar to the DSF Gated Power Sequence, but the Box receiving the power up command powers itself before sending the CMI_POWER_ENABLE and CMI_SSTART message to DSF ports.**
- **The message with CMI_WAKE and CMI_CLP_READY is not sent until messages are received from all DSF ports**

Cable Management Power- MGTPWR

- **Previous revisions of the specification called for 8 mA of MGTPWR to be supplied to the cable.**
- **Analysis has shown that the need may be higher. SAS cables in particular have changed to require 30 mA of current for microcontrollers.**
- **The 0.7 draft increase current to MGTPWR to 30 mA.**

Cable Skew

- **Cable aggregation allows the use of multiple cable assemblies to create a single link.**
- **The skew allowed for the cable assembly is 2 ns.**
- **For aggregated assemblies, this is 2 ns across all assemblies for the link.**

- **This creates a need for a mechanism to determine if cables are compatible to be combined.**

External Cable Skew Budget

Jitter Parameters	Symbol	Values (RMS in ps)	Notes
Total_Interconnect_Skew	S_T	2.7 ns	1, 2, 3
Subsystem Skew	S_S	0.35 ns	4
Cable Assembly Skew	S_C	2.0 ns	5

Notes:

1. This does not include Transmitter output skew.
2. This is the total skew allowed for a link. If multiple cable assemblies are used, the flight time delay information (Table 6-3 bytes 108-109) must be read from all aggregated cable assemblies in the link and must meet the Total Interconnect Skew for the entire link.
3. The skew at any point is measured at the zero crossings of differential voltage of the compliance pattern, while simultaneously transmitting on all lanes. The compliance pattern is defined by the PCI Express Base Specification.
4. The limit is based on approximately two inches of stripline.
5. This limit is across all signal pairs. When using multiple cable assemblies for a port aggregation, the total skew across all pairs in the link must meet this requirement.

Reporting Cable Delays

- **The 0.6 draft had a 1 ns resolution for reporting delay**
  - We used the ceiling of the measured delay of the shorted diff pair (6.2.2.3.2)
- **A skew of 1ns is specified for a single cable assembly** (Table 3-2)
- **With 1ns of resolution and 1ns of skew, cable assemblies must match for aggregation** (Table 3-2 Note 8)

Skew Compatibility with 1ns Limits

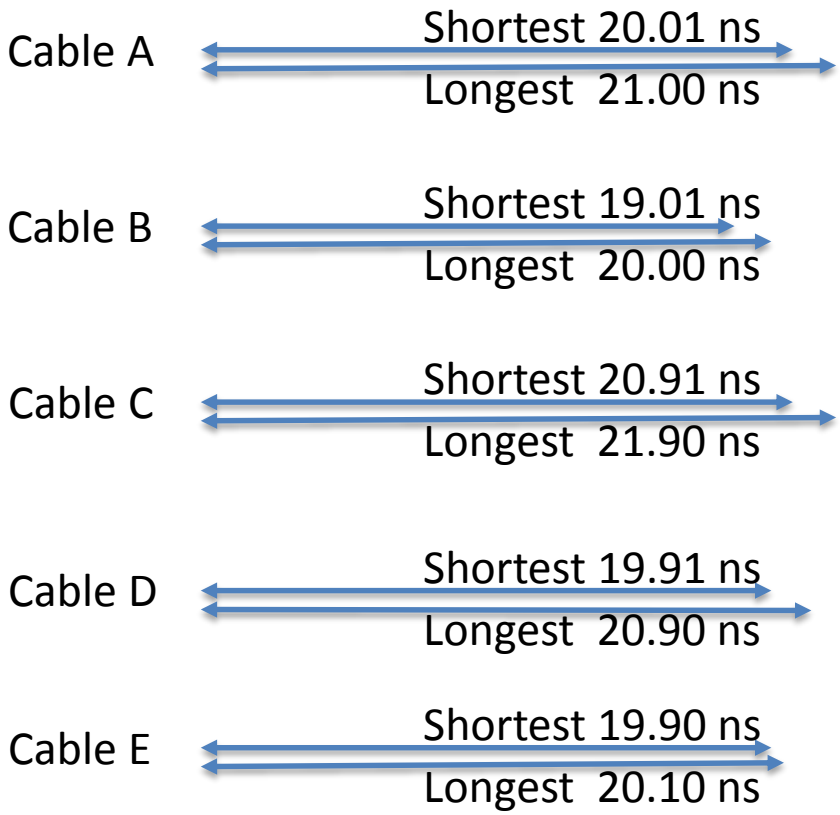

Reported Length (ns)*	Specification Compatibility	Functional Capability (<2ns total skew)
21	C	B (1.99 ns), C (1.89 ns), D (1.09 ns), E (1.1 ns)
20	D,E	A (1.99 ns), D (1.89 ns), E (1.09 ns)
21	A	A (1.89 ns), D (1.99 ns), E(2.0 ns)
20	B, E	A (1.09 ns), B (1.89 ns), C (1.99 ns), E (1.0 ns)
20	B,D	A (1.1 ns), B (1.09 ns), C(2.0 ns), D (1.0 ns)

Resolution for Skew During Cable Aggregation

- **Repurpose the sixteen bits originally designated for reporting Propagation Delay in 1 ns increments.**
  - Use thirteen bits to report up to 2 us of delay with a 250ps resolution. This information is used to create the *Base* parameter.
  - Use three bits to report the skew within the cable assembly in 250 ps of resolution. This information is used to create the *Span* parameter.

How this Works for Cable Assembly Providers

- **Lower Bound**
  - The fastest measured propagation delay of the pairs within a cable assembly rounded down to the nearest multiple of 250ps.
- **Upper Bound**
  - The slowest measured propagation delay of the pairs within a cable assembly rounded up to the nearest multiple of 250ps.

  Note: this creates up to a 500ps skew within the reported cable assembly.

- **Base Parameter**
  - Base = Lower Bound / 250 ps
- **Span Parameter**
  - Span = (((Upper Bound – Lower Bound)/250 ps) – 2) for UB-LB > 250ps
  - Span = 0 for UB - LB ≤ 250 ps

How this Works for Cable Management Interface Controllers Attempting to Aggregate Cables Assemblies

- **The CMI controller imports the Base and Span parameters from the cable assemblies for the assemblies attached to the ports to be aggregated.**

- **The CMI controller recalculates the Upper and Lower Bound for each assembly and determines if the skew is under 2ns.**
  - Lower Bound = Base * 250 ps
  - Upper Bound = (Base + Span + 2) * 250 ps
  - The smallest base among the cables to aggregate is the reference to compare the upper bounds of all the other cables.

Cable Aggregation Calculation Examples for Base and Span

Parameter	Cable 1	Cable 2	Cable 3	Cable 4
Fastest Propagation Delay	10.2 ns	133.6 ns	794.8 ns	2047.9 ns
Slowest Propagation Delay	12.1 ns	133.7 ns	795.2 ns	2049.2 ns
Lower Bound	10.00 ns	133.50 ns	794.75 ns	2047.75 ns
Upper Bound	12.25 ns	133.75 ns	795.25 ns	2049.25 ns
Upper Bound – Lower Bound	2.25 ns	0.25 ns	0.50 ns	1.5 ns
Base	40d (28h)	534d (216h)	3197d (C6Bh)	8191d (1FFFh)
Span	7h (111b)	0h (000b)	0h (000b)	4h (100b)

CMI Controller compares Upper Bounds to Lower Bounds to ensure skew is less than 2ns.

Note: Cable suppliers are to ensure cables are under 2ns in skew (Table 3-2) so the largest possible reported Span valued of 7h (or 2.25ns) is acceptable and does not reflect “real” skew.

Call to Action

- **The 0.9 draft is in various stages of review.**

- **When notified, take time to review the specification.**

Thank you for attending the PCI-SIG Developers Conference 2017.

For more information, please go to www.pcisig.com

Don't forget to submit your feedback via the mobile app!

Download the app by searching stores for **PCI-SIG Developers Conference** or entering the following URL into your mobile browser: **https://crowd.cc/s/EOtK**

Enter event code: **DevCon17**

Alternatively, access here: https://event.crowdcompass.com/devcon-17

Note: Create an account within the app so Admin knows who to contact if selected as the prize winner.

You will earn 1 raffle entry for each session you provide feedback on (up to 11 sessions).
General survey feedback = 1 raffle entry.